PHOTONIC A/D CONVERTERS FOR WIDEBAND AND HIGH-DYNAMIC-RANGE PERFORMANCE

J.C. Twichell, Z.J. Lemnios
Lincoln Laboratory, Massachusetts Institute of Technology
Lexington, MA 02420-9108
C.E. Dickerson
Program Executive Office
Surface Combatants/AEGIS Program
Arlington, VA 22202

Abstract

Navy TBMD AEGIS ships will be required to perform radar tracking and exoatmospheric discrimination against future TBM threats that will have significantly greater range, higher velocity, and lower radar cross section than the current predominate threat. A significant contributor to current radar tracking and discrimination limitations is the downconversion chain to intermediate frequency (IF) and the analog-to-digital (A/D) conversion. Conventional multistage downconversion receivers are limited in dynamic range, linearity, and A/D conversion rate. Performing the A/D conversion at RF rather than IF can provide a solution to these limitations as well as a significant increase in radar sensitivity. Technology development for photonic A/D converters could lead to a new generation of receivers with superior linear dynamic range and sensitivity to counter the future TBM threat.

This paper presents a novel approach to A/D conversion based upon precise optical sampling and optical phase discrimination, with the proposed development of a low jitter mode-locked laser (< 50 fs), and the development of charge-based quantizers with greater than 14 bit linearity (at 6 gigasamples per second). Initial measurements will be presented for UHF and S-band architectures supporting the extension of A/D converter performance well beyond the commercial regime. Converters based on this technology would enable a new class of digital receivers with about 12 dB improvement in signal-to-noise ratio, a 90 dB spurious-free dynamic range, and a 100-fold improvement in instantaneous bandwidth over conventional receiver approaches.

Introduction

Analog-to-digital (A/D) converters limit the amount and accuracy of data which can be collected with a radar system. As more challenging threats emerge, systems must be developed with greater resolution and higher bandwidth. In broad terms, converters with SNR > 70 dB fs, and third-order intermodulation (IM3) spurious-free dynamic range (SFDR) in the 90 dBc regime at speeds of a gigasample per second (GS/s) or higher are needed. Achieving these objectives simultaneously is beyond the reach of conventional electronics. Performing critical elements of the converter function in the optical domain provides a dramatic improvement in both speed and precision.

Analog-to-digital converters typically comprise a sample-and-hold function and a quantizer. The sample-and-hold element represents the time dependent input signal at a precise, triggered time. This value is held constant at the output of the sample-and-hold circuit and provides a fixed input value to the quantizer. The quantizer takes this input and converts the fixed sample to a numeric value.

Errors enter this signal path in several ways. The sample and hold must have a very fast amplifier which drives an unreasonable load, a capacitor. Often nonlinearities limit the performance of this amplifier. Timing errors result in sampling the input signal at a time different from the expected time. This contributes an error which depends on the time-rate-of-change of the input signal. There are two components to timing errors. Jitter, or a stochastic timing error, contributes stochastic noise. Timing variations which are a function of the input signal produce errors much like the nonlinear response of the input amplifier. Finally, the quantizer also makes errors. These include differential (value to value) and integral (full-scale spanning) nonlinearities, but

quantizers also have dynamic errors, and smaller scale integral nonlinearities as well. Typically, the errors in the sample and hold dominate, limiting the performance of the converter.

Three primary measures are used to describe the performance of A/D converters. The signal-to-noise ration (SNR) measures stochastic noise. SFDR indicates the residual harmonic content induced by the nonideal character of the measurement. Finally, IM3 is the third-order intermodulation distortion. The stochastic noise can be reduced by averaging in a signal processor. In narrowband systems, spurs can often be placed in benign locations by careful choice of sampling frequency. Wideband systems do not have this flexibility. Odd intermodulation orders are always in band. For radars and communication systems IM3 performance is often the most difficult requirement to meet.

Figure 1, courtesy of R. Walden at Hughes Research, shows the SNR versus sampling rate of present-day converters. It should be stressed that the detailed mechanism responsible for the limiting SNR likely is different in different speed regimes. For example, at high speeds, noise in the ground reference may well contribute substantially to the SNR. The DARPA A/D program built the device at 3 GHz and 40 dB. Commercial devices recently have advanced rapidly towards the "Walden wall." The Walden wall is moving upwards at the rate of about 1 bit every eight years. It is our contention that by shifting to the optical domain, the limitations of electrical devices can be circumvented.

Many applications are well beyond the Walden wall. These include modest extensions to narrowband radars requiring several 10's of megahertz of bandwidth, cellular base stations requiring 75 MHz of bandwidth, and wideband applications. Imaging radar and wideband communications require bandwidths in excess of 1 GHz.

Tests to date have explored 0.5 and 10 MHz bandwidth (1 and 20 MS/s) at carrier frequencies of 500 MHz, and 3 GHz. Programs in place are aimed at bandwidths of up to 250 MHz (500 MS/s), with long-term goals of 3 GHz (6 GS/s). Sampling in the optical domain is the fundamental technique which permits this advance.

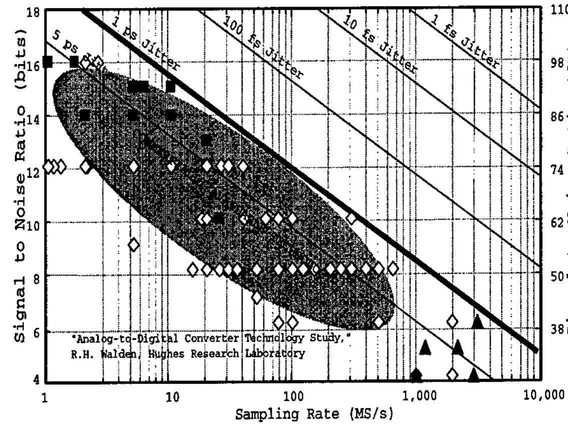

Figure 1. The signal to noise ratio (SNR) versus sampling speed for a wide array of analog-to-digital (A/D) converters. The data reflects 1997 parts. The diagonal lines represent the SNR which would be a consequence of a given level of sampling jitter for Nyquist sampling versus converter speed. No commercial device has an effective sampling jitter of less than 1 ps, which has been dubbed the "Walden wall."

Optical Sampling

The basic approach to optical sampling is not new. Figure 2 shows a diagram of a basic optical sampler. Mode-locked lasers offer extraordinary precision in the timing of the train of output pulses. These can be used to sample an electrical signal applied to an electro-optic modulator, producing a sampled representation of the input electrical signal both in the optical domain and in the electrical domain after the pulses are detected. The problem with this approach is the linearity with which an electrical signal can produce an amplitude-modulated optical signal.

Outstanding linearity can be achieved by approaching the interferometer as a phase modulator and the front end of a phase demodulator, as shown in figure 3. The interferometer is shown in simplified form. The applied voltage produces an optical phase shift as shown in the plot on the left. When mixed with the reference arm of the interferometer, classic fringes result. These fringes are sinusoidal because, to very good accuracy, the light in the interferometer can be represented as sinusoidal. The transfer function of the

interferometer can be simply inverted. The plot on the right shows the inferred phase, hence (to within a constant) the input voltage, resulting from this inversion.

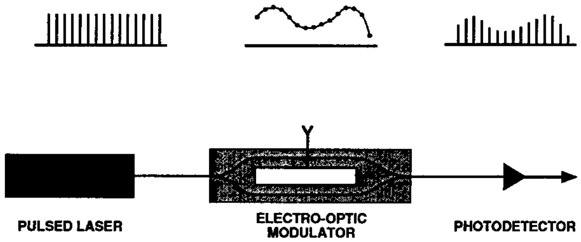

Figure 2. Conceptual diagram of an optical sampler. A laser generates a train of short, precisely timed optical pulses. These pulses are modulated by an electrical signal in an electro-optic modulator. The result is a sampled representation of the signal. Note that after the modulator, the required timing precision is reduced to the interpulse interval, rather than the femtosecond precision needed going into the modulator.

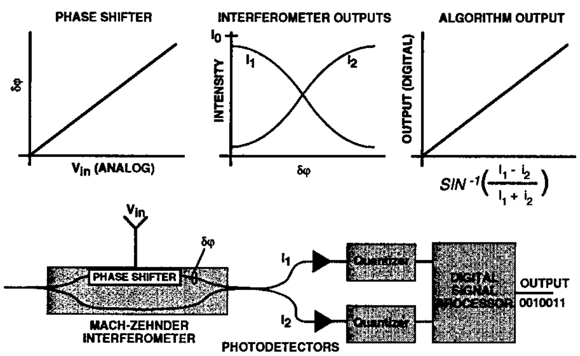

Figure 3. Phase encoding a signal in the optical domain. See the text for detailed discussion.

Several important aspects of this approach should be noted. First, as one would expect for a phase demodulator, the inference of the applied voltage is insensitive to the optical carrier amplitude. The mechanism is clearly seen in the denominator of the inverse transfer function which is simply the total detected light. Since the interferometer conserves photons (except for constant losses), this approach should provide independence to amplitude noise in the laser. Indeed, we have measured 60 dB attenuation of laser noise in the inferred phase, and even this was likely due to chirp due to the strong modulation of the laser.

A more subtle consequence of this amplitude independence is that the quantizers can be dithered without putting any signal in-band. Consider a constant input signal. If the laser pulse amplitude varies, the quantizers will measure varying values, but the inferred signal will not change. Thus, the quantizers explore a range of codes for a given input signal. This converts coherent errors in the quantizer, the primary cause of spurs, to noise allowing them to be averaged away. In a similar vein, the technique is insensitive to nonidealities in the optical paths, since the signal is encoded. Thus, the finite contrast of the interferometer (seen in the center plot) causes a tiny degradation in SNR, but has no impact on accuracy or linearity. The data shown later has a 10% gain error and about 1% offset error deliberately left in place, again with only trivial consequences. Finally, note that the approach gives optical isolation of the timing reference (the mode-locked laser), the input signal, and the quantizer and digital signal processor.

The fundamental advantages optics has over electronics for timing stability are low dispersion and low fractional bandwidth. The optical cavity of a mode-locked laser provides gain in the optical domain where the fractional bandwidth is tiny and the dispersion is also very small. Formally, the edge rates used to establish the sampling time (and the sensitivity to noise) are limited by dispersion in the transmission lines, either optical or electrical. Clearly, this represents a rather fundamental limitation for the dielectrics of integrated circuits.

Sampling in the optical domain also offers a significant advantage for the linearity of the sampling process. The sampling signal simply does not interact with the input electrical signal in an optical sampler. In the electrical domain, this interaction is always present, producing timing errors which vary coherently with the input signal.

Experimental Program

An experimental program began with a small, conceptual test. This allowed us to explore the basic concept and address what was originally considered the primary risk, the linearity of the phase modulator. To our astonishment, we were unable to measure the nonlinearity given the noise floor of our equipment. An extension to our original tests has been constructed. The objective of the Technology Extension project was to push an order of magnitude simultaneously in four directions: The sampling speed was increased from 500 MHz to 3 GHz, the

bandwidth increased from 0.5 to 10 MHz, the SFDR increased from 78 to 90 dB, and the IM3 from 90 to > 100 dB. Initial results have been very encouraging. A DARPA program is in place to push the bandwidth by multiplexing quantizers. Because the signal is encoded, one gains relief from the mechanisms which normally make the multiplexed approach intractable. We will return to this subject later.

The initial setup of the conceptual test experiment was quite simple. Figure 4 shows the configuration. The object of the experiment was *not* to explore low-jitter sampling. Rather, the intent was to test the limits of the phase modulation approach without regard to bandwidth at a relevant carrier frequency. Figure 5 shows the timing of the various components. While the laser is running at 500 MHz, the two-tone test source was operated very close to this frequency. This aliased the test signal to baseband, and let us use low bandwidth digitizers to characterize the linearity of the system. The results are shown in figures 6 and 7. In 0.5 MHz of bandwidth at 500 MHz, this test demonstrated 90 dB IM3 SFDR and 78 dB SFDR.

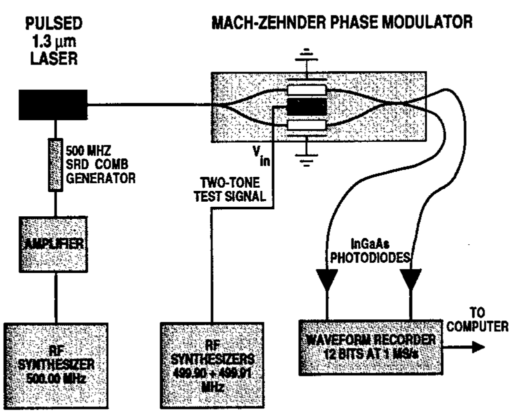

Figure 4. The initial test setup for characterizing the linearity of an optical sampler. A 1.3 µm Fabry-Perot laser was gain switched. The laser was driven from a 500 MHz synthesizer and a 1 W amplifier which drove a step-recovery diode. The result was a train of ~100 ps pulses at 500 MHz. These were used to sample a two-tone test signal in a commercial Y-branched balance modulator (YBBM). The two tones were extremely clean, generated 100 kHz from the laser frequency, and 10 kHz apart. The output of the interferometer was detected and digitized with an antique 12 bit Lecroy digitizer.

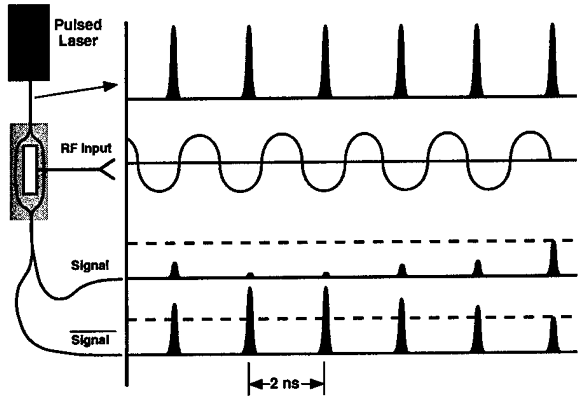

Figure 5. This figure illustrates the operation of the experiment. A train of pulses from the laser enters the modulator. The applied RF is close to the laser frequency (here the frequency difference is greatly exaggerated). The output of the interferometer is a pair of complementary pulse trains. The envelope of these pulses is the down-sampled signal at the difference between the laser repetition rate and the RF signal. In the experiment, about 100 laser pulses were averaged per measurement.

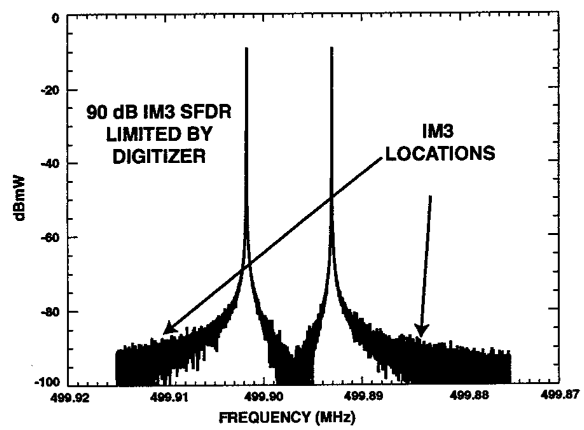

Figure 6. The third order intermodulation (IM3) distortion should show up as two peaks at the indicated locations. They are not visible in the noise, which is 84 dB below the single-tone amplitude. The A/D community uses the coherent sum of the two tones as the reference amplitude, adding 6 dB to the signal amplitude. Thus the intermodulation distortion is 90 dB below the coherent sum of the tones.

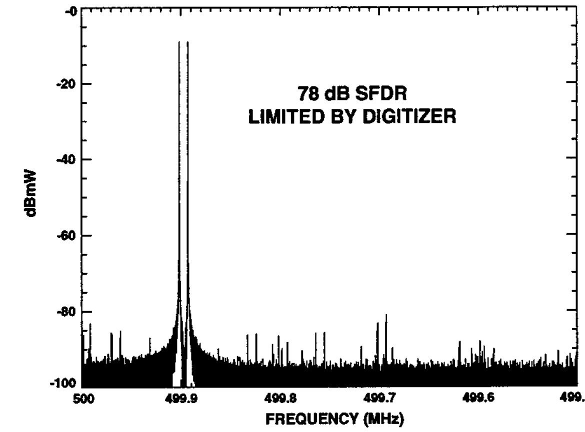

Figure 7. The full 0.5 MHz bandwidth of the conceptual test experiment. The spurious free dynamic range (SFDR) was 78 dB, limited by the digitizer used.

Figure 8 shows the impact of phase encoding. The result is an SFDR of 57 dB, limited now by the harmonic spurs and IM3. Note the strength of the third-order error. When perfectly biased at quadrature, the second-order error is zero and only the third-order error remains. The phase discrimination also removes laser amplitude noise. Over 60 dB of amplitude suppression has been demonstrated.

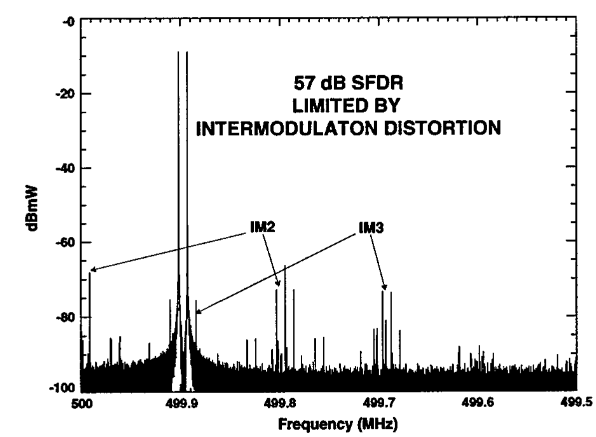

Figure 8. The SFDR treating the system as a normalized amplitude detector. The signal plotted is the difference over the sum of the two detectors. This procedure removes the amplitude variations in the laser, but treats the optical signal as simple differential amplitude modulation.

Tests were performed to assess the sensitivity of the linearity to a number of errors in the system. It was found that 1% precision of the gain, offset, and contrast ratio was more than adequate to maintain 90 dB of linearity. Detector offsets of 10 mV are tolerable. The detector quantum efficiency, and differential optical transmission must be calibrated to within 1%. Both these requirements are easily met over temperature and supply variations with ordinary commercial off-the-shelf (COTS) components. The interferometer extinction ratio must be known to within 2% to maintain 90 dB of linearity. Vendors typically guarantee the extinction ratio will be better than 20 dB (i.e., extinction > 0.99) with typical values of 30 dB (0.999). The impact of treating the signals as phase encoded removes most of the concern about amplitude errors in the signal paths.

The Technology Extension experiment shown in figure 9 was built to push the bandwidth, operating frequency, and dynamic range. The laser wavelength was shifted to 1.55 μm to take advantage of commercial telecommunications industry technology, for example, erbium fiber amplifiers. For the data presented below, the laser was operated at 1 GHz, producing 16 to 22 ps pulses. The test tones were offset from the third harmonic of the laser repetition rate. This enhanced the linearity errors by a factor of 3.

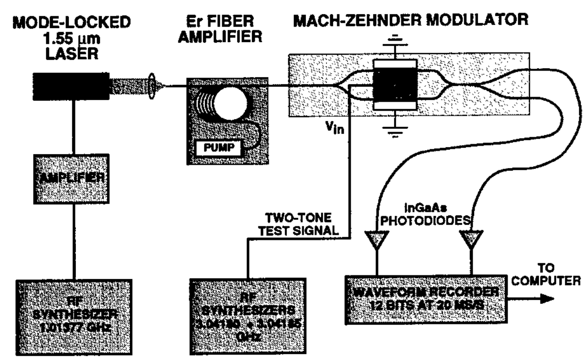

Figure 9. The configuration of the Technology Extension experiment. A mode-locked 1.55 μm semiconductor laser operates at 1 or 3 GHz. An erbium fiber amplifier is used to increase the amplitude of the pulses. Not shown is an optical fiber to remove pump and background ASE. the digitizers used at HP1437A with >90 dB SFDR. The test signal is two tones 50 kHz apart offset from the third harmonic of the laser frequency by 500 kHz.

The laser pulse was characterized in the time domain with a 50 GHz sampling oscilloscope and a 45 GHz detector. The measured pulsewidth was about 11 ps after deconvolving the instrument response. A small tail was noticed trailing the pulse.

This tail is believed to be the source of the limited IM3 distortion observed. The laser is mode locked by gain modulation of the entire gain region of the laser. This leaves a substantial excited population after the gain drops below 1 and the pulse terminates.

The results of the initial testing shown in figure 10, indicate good performance at 3 GHz. The bandwidth of the system was increased to 10 MHz using 20 MS/s digitizers. The intermodulation distortion was measured as 81 dB limited, we believe, by background light from the laser. The SFDR was 72 dB limited by imperfect shielding of the photodetectors.

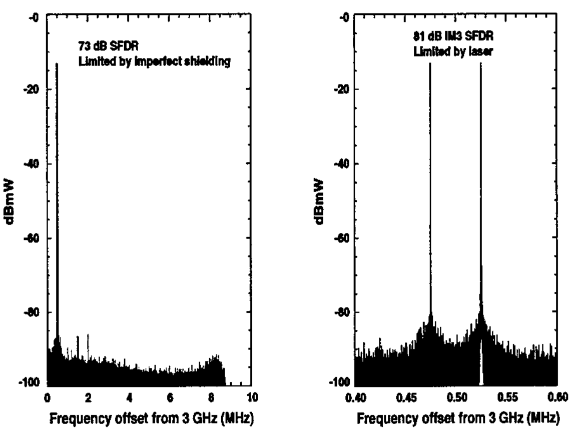

Figure 10. Technology Extension experiment results. The SFDR was 73 dB, limited by imperfect shielding. The IM3 SFDR was 81 dB. The intermodulation products are just barely visible. The tone amplitude was adjusted for maximum dynamic range.

It should be noted that with appropriate front end filters to limit the input bandwidth, both these experimental tests would make outstanding receivers. The simplicity of the approach limits the number of components in the signal path. The two experiments demonstrate that optical sampling can be made extraordinarily linear, and that this linearity is maintained at microwave frequencies.

Analysis

Shot noise limits the performance of any optical system carrying information. There must be enough photons within the bandwidth of the measurement to provide the desired resolution. The approach we have taken faces additional constraints. The maximum average power which can be coupled into a monolithic electro-optic modulator is less than 1 W. This limits the pulse energy which in turn limits our ability to resolve the energy deposited in the detector.

The modulation depth is the ratio of a full-scale change in signal to the average amplitude of the signal. As the modulation depth is increased, the impact of amplitude errors in phase demodulation becomes more severe. At the same time, however, we need fewer photons to resolve the full-scale change in intensity.

Figure 11 shows the laser energy per pulse required as a function of modulation depth. The power limit for the fiber to $LiNbO_3$ coupling is an average power constraint, and is optimistically 1 W. Vendors typically will only guarantee 200 mW. For a given pulse rate (the measurement rate) the average power limit imposes a limit on the energy per measurement. This, in turn imposes a constraint on the minimum modulation depth which can be used. As the modulation depth increases, fewer photons are needed to resolve 1 part in 4096 (12 bits) of the modulation depth used. For example, at 20% modulation depth, the intensity must be resolved to 1 part in 5×4096, requiring 25 times as many photons as for a 100% modulation depth. The practical upper limit to modulation depth for a two-port interferometer is about 50%.

The photodetectors integrate the optical signal over the pulse. Thus our measure of the encoded phase is an average over some short time interval. Unfortunately, averaging the intensity in time is *not* the same as averaging the signal in time. As the modulation depth increases, the error made by this time integration increases. This error shows up as third order intermodulation distortion. There is, then, a limit on the width of the laser pulse relative to the highest input frequency, and this limit is a function of the modulation depth used. This set of trades is illustrated in figure 12. For the modulation depths required by the shot noise constraints it is clear that the pulsewidth must be restricted to only a few percent of the highest frequency measured. At S-band this requires pulse widths less than 10 ps, well within the capabilities of COTS lasers.

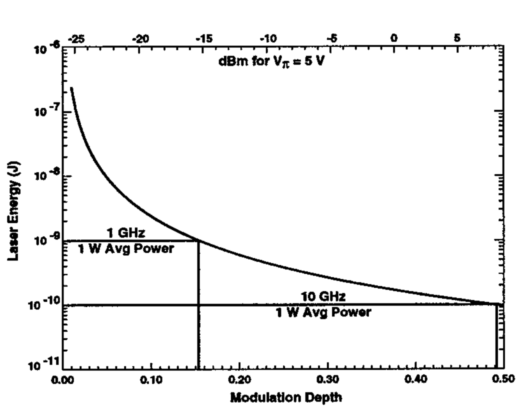

Figure 11. Shot noise limits for optical sampling as a function of modulation depth. The vertical axis is the energy required per measurement to 12 bits of precision for each modulation depth. The top shows the power in dBm for V_π of 5 V. The minimum modulation depth given the 1 W average power limitation for two cases of interest are shown.

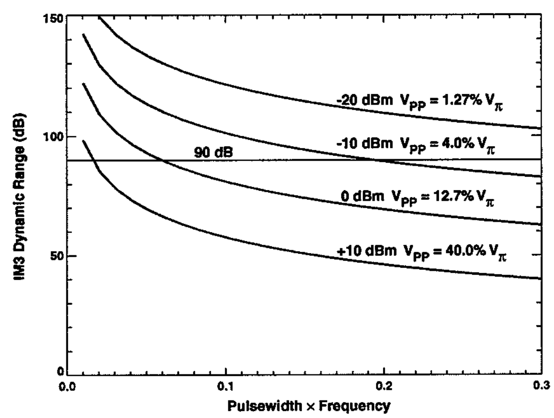

Figure 12. Third-order intermodulation distortion is plotted as a function of fraction of the RF tone period for several modulation depths. For 3 GHz RF, pulsewidths must be held to less than 10 ps.

A similar class of errors occurs if there is a constant level of illumination in addition to the pulse. This background illumination could be spontaneous emission from the laser or pump light leaking through the filters of the laser. The photodetectors integrate this light over their staring interval. This leaves a component of the detected signal as a long-time average of the phase-encoded input signal. Again, the time average of the phase-encoded signal is not the average of the input signal. We therefore expect errors dependent on the strength of the optical DC illumination and the modulation depth used by the modulator. A model was constructed to illustrate the severity of this process. It was found that for large modulation depths the background light from the laser must be held to nearly undetectable limits. This presents a severe challenge for semiconductor mode-locked lasers. Er fiber lasers, with their much longer excited state lifetimes, will have much lower spontaneous emission. The pump light must be well filtered from the output if this advantage is to be realized.

Wideband Extensions

Sampling in the optical domain has been demonstrated at high speed with good linearity. The bandwidth of our measurements are limited by the quantizer. The primary issue now becomes: How do we digitize faster?

The obvious solution is to build a faster quantizer. The design of this faster quantizer is free of the constraints for the track and hold. This may provide a factor of 2 to 4 improvement in the speed of the quantizer. Given the speed of the optical sampler, we clearly need a better solution. It would appear we have no choice but to use multiple quantizers to provide the bandwidth desired. We contend it will be easier to equalize multiple quantizers measuring phase-encoded information than it will be to build a single extremely fast quantizer.

We have two choices of signal demultiplexing technique to use multiple quantizers. Electrically demultiplexing is simple, but it is difficult to maintain linearity at high speed. Optical demultiplexing provides outstanding linearity at the expense of complexity. We will employ a combination of these techniques. Optical demultiplexing will be used until the signal bandwidth is low enough that the electrical demultiplexers can maintain the required linearity. One of the objectives of our present effort is to establish just where this boundary is.

The baseline architecture is shown in figure 13. An advanced mode-locked laser generates a 6 GHz train of 10 ps pulses. The input electrical signal is applied to the interferometer, both legs of which are subsequently digitized. The pulses are demultiplexed in the optical domain to reduce the bandwidth at each output port. The pulses are detected and further demultiplexed to match the bandwidth of the

quantizers. Ultimately, the speed of an individual quantizer establishes how much multiplexing must be done. A linearity of the electrical demultiplexer will set how much of this demultiplexing must be done in the optical domain.

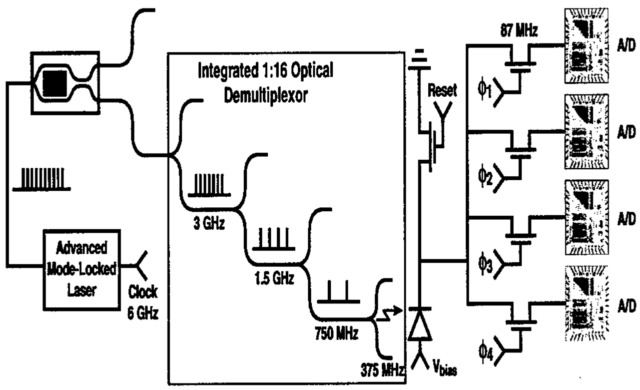

Figure 13. Architecture of a wideband converter. The bandwidth of the sampled signal is reduced in the optical demultiplexer until further processing can be done in the electrical domain.

The charge domain converter (CDC) offers a near ideal match to the needs of the optical sampling systems. In this device the signal is represented as a packet of charge, not a voltage. The photodetector in our optical sampler detects a pulse of light which directly generates a packet of charge. We then go to great lengths to convert it to a voltage linearly. This is unnecessary for the CDC. Recent results for the charge domain digitizer indicate 15 bits of differential linearity and SNR in excess of 90 dB can be achieved with a power budget of only 18 mW/MHz.

Several constraints must be addressed before optical sampling can be considered a viable technology. The timing stability of compact, efficient mode-locked lasers must be improved by about a factor of 5. A better quantizer (smaller, faster, and lower power) is clearly needed. Maintaining linearity of the photodetectors, and the transfer of the signal from the detector to the digitizer, will always be a challenge. Finally, note that optical sampling may well eliminate the A/D converter as the technical constraint in many systems. The timing stability of the reference oscillator will limit performance for some systems.

In summary, optical sampling offers a path to converter performance which has been inaccessible to date. For systems in the few hundred MS/s regime optical sampling offers order-of-magnitude improvements in linearity and complexity. There is a clear path to extend this technology to wideband applications.

Acknowledgments

This work was sponsored by the Defense Advanced Research Projects Agency.